

LA-UR 95-1504

Approved for public release; distribution is unlimited

Approximation Algorithms for PSPACE-Hard Hierarchically and Periodically-specified Problems

Authors: M.V. Marathe, H.B. Hunt III, R.E. Stearns,
V. Radhakrishnan

LOS ALAMOS

NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse this viewpoint of a publication or guarantee its technical correctness.

Approximation Algorithms for PSPACE-Hard Hierarchically and Periodically Specified Problems

Madhav V. Marathe³ Harry B. Hunt III¹ Richard E. Stearns¹ Venkatesh Radhakrishnan²

Abstract

We study the efficient approximability of basic graph and logic problems in the literature when instances are specified hierarchically as in [Le89] or are specified by 1-dimensional finite narrow periodic specifications as in [Wa93]. We show that, for most of the problems Π considered when specified using **k-level-restricted** hierarchical specifications or k -narrow periodic specifications, the following holds:

- Let ρ be any performance guarantee of a polynomial time approximation algorithm for Π , when instances are specified using standard specifications. Then $\forall \epsilon > 0$, Π has a polynomial time approximation algorithm with performance guarantee $(1 + \epsilon)\rho$.
- Π has a polynomial time approximation scheme when restricted to planar instances.

These are the first polynomial time approximation schemes for PSPACE-hard hierarchically or periodically specified problems. Since several of the problems considered are PSPACE-hard, our results provide the first examples of natural PSPACE-hard optimization problems that have polynomial time approximation schemes. This answers an open question in Condon et. al. [CF+93].

Keywords. hierarchical specifications, periodic specifications, PSPACE-hardness, approximation algorithms, computational complexity, CAD systems, VLSI design.

AMS(MOS) subject classification. 68R10, 68Q15, 68Q25, 05C40.

¹Email addresses: {hunt,res}@cs.albany.edu. Department of Computer Science, University at Albany - SUNY, Albany, NY 12222. Supported by NSF Grants CCR 89-03319 and CCR 94-06611.

²Mailstop 47LA, Hewlett-Packard Company, 19447 Pruneridge Avenue, Cupertino, California 95014-9913. Email: rven@cup.hp.com

³Part of the research was done when the author was at SUNY-Albany, and was supported by NSF Grant CCR 94-06611. Current address: P.O. Box 1663, MS K990, Los Alamos National Laboratory, Los Alamos NM 87545. Email: madhav@c3.lanl.gov. The work is supported by the Department of Energy under Contract W-7405-ENG-36.

⁴A preliminary version of this paper appeared in the *Proc. 26th ACM Annual Symposium on Theory of Computing (STOC)*, 1994, pp. 468-477.

1 Introduction and motivation

Many practical applications of graph theory and combinatorial optimization in CAD systems, mechanical engineering, VLSI design and software engineering involve processing large objects constructed in a systematic manner from smaller and more manageable components. An important example of this occurs in VLSI technology. Currently, VLSI circuits can consist of millions of transistors. But such large circuits usually have a highly regular design and consequently are defined systematically, in terms of smaller circuits. As a result, the graphs that abstract the structure and operation of the underlying circuits (designs) also have a regular structure and are defined systematically in terms of smaller graphs. Methods for describing large but regular objects by small descriptions are referred to as *succinct specifications*. Over the last twenty years several theoretical models have been put forward to succinctly represent objects such as graphs and circuits. (see for example [BOW83, CM91, Ga82, HW92, IS87, KMW67, KO91, Le88, Or82a, Or84b, Wa84]). Here, we study two kinds of succinct specifications, namely, hierarchical and periodic specifications.

Hierarchical specifications allow the overall design of an object to be partitioned into the design of a collection of modules; which is a much more manageable task than producing a complete design in one step. Such a top down (or hierarchical design) approach also facilitates the development of computer aided design (CAD) systems, since low-level objects can be incorporated into libraries and can thus be made available as submodules to designers of large scale objects. Other areas where hierarchical specifications have found applications are VLSI design and layout [HLW92, HW92, RH93], finite element analysis, software engineering and datalog queries (see [HLW92, Ma94] and the references therein). Periodic specifications can also be used to define large scale systems with highly regular structures. Using periodic specifications, large objects are described as repetitive connections of a basic module. Frequently, the modules are connected in a linear fashion, but the basic modules can also be repeated in two or higher dimensional patterns. Periodic specifications are also used to model time variant problems, where the constraints or demands for any one period is the same as those for preceding or succeeding periods. Periodic specifications have applications in such diverse areas as transportation planning [Or82a, HLW92, Ma94], parallel programming [HLW92, KMW67] and VLSI design [IS87, IS88].

Typically, the kinds of hierarchical and periodic specifications studied in the literature are generalizations of standard specifications used to describe objects. An important feature of both these kinds of specifications is that they can be much more concise in describing objects than standard specifications. In particular, the size of an object can be exponential in the size of its periodic or hierarchical specifications. As a result of this, problems for hierarchically and periodically specified inputs often become PSPACE-hard, NEXPTIME-hard, etc.

In this paper, we concentrate our attention on (i) the hierarchical specifications of Lengauer [Le86, Le88, Le89] (referred to as L-specifications) and (ii) the 1-dimensional finite periodic specifications of Gale and Wanke [Ga59, Wa93] (referred to as 1-FPN-specifications). Both of these specifications have been used to model problems in areas such as CAD systems and VLSI

design [Le89, LW87a, Le90], transportation planning [Ga59], parallel programming [Wa93], etc. We give formal definitions of these specifications in Sections 4 and 5.

Let Π be a problem posed for instances specified using *standard* specifications. For example, if Π is a satisfiability problem for CNF formulas, the standard specification is sets of clauses, with each clause being a set of literals. Similarly if Π is a graph problem, the adjacency matrix representation or the adjacency list representation of the edges in the graph are standard specifications. For the rest of the paper, we use

1. L- Π to denote the problem Π , when instances are specified using the hierarchical specifications of Lengauer [Le89] (see Definition 4.1), and
2. 1-FPN- Π to denote the problem Π , when instances are specified using the 1-dimensional finite periodic specifications of Wanke [Wa93] (see Definition 5.1).

Thus for example, L-3SAT denotes the problem 3SAT when instances are specified using L-specifications and 1-FPN-3SAT denotes the problem 3SAT when instances are specified using 1-FPN-specifications. For the rest of this paper, we use the term *succinct specifications* to mean both L-specifications and 1-FPN-specifications.

2 Summary of results

In this paper, we discuss a natural syntactic restriction on the L-specifications and call the resulting specifications **level-restricted** specifications. (For 1-FPN-specifications our notion of level restricted specifications closely coincides with Orlin’s notion of narrow specifications [Or82a].) Most of the problems considered in this paper are PSPACE-hard even for level restricted specifications (see [Or82a, LW92, MH+95a]). Consequently, we focus our attention on devising polynomial time approximation algorithms for level restricted L- or 1-FPN-specified problems. Recall that an approximation algorithm for a minimization problem⁵ Π provides a **performance guarantee** of ρ if for every instance I of Π , the solution value returned by the approximation algorithm is within a factor ρ of the optimal value for I . A **polynomial time approximation scheme** (PTAS) for problem Π is a family of algorithms such that, given an instance I of Π and an $\epsilon > 0$, there is a polynomial time algorithm in the family that returns a solution which is within a factor $(1 + \epsilon)$ of the optimal value for I . The main contributions of this paper include the following.

- We design polynomial time approximation algorithms (for arbitrary instances) and approximation schemes (for planar instances) for a variety of natural PSPACE-hard problems specified using either level restricted L- or 1-FPN-specifications. These are the first polynomial time approximation schemes in the literature for “hard” problems specified using either L- or 1-FPN-specifications.

⁵A similar definition can be given for maximization problems.

To obtain our results we devise a new technique called the *partial expansion*. The technique has two desirable features. First, it works for a large class of problems and second, it works well for *both* L-specified and 1-FPN-specified problems.

- For problems specified using level-restricted L- or 1-FPN-specifications, we devise polynomial time approximation algorithms with performance guarantees that are asymptotically equal to the best possible performance guarantees for the corresponding problems specified using standard specifications.
- These results presented in this paper are a step towards finding sufficient syntactic restrictions on the L- or 1-FPN-specifications that allow us to specify a number of realistic designs in a succinct manner while making them amenable for rapid processing.

Our results provide the first examples of *natural* PSPACE-complete problems whose optimization versions have polynomial time approximation schemes. Thus they affirmatively answer the question posed by Condon, Feigenbaum, Lund and Shor [CF+93] of whether there exist *natural* classes of PSPACE-hard optimization problems that have polynomial time approximation schemes.

2.1 Meaning of approximation algorithms for succinctly specified problems

When objects are represented using L- or 1-FPN-specifications, there are several possible ways of defining what it means to “design a polynomial time approximation algorithm”. Corresponding to each decision problem Π , specified using either L- or 1-FPN-specifications, we consider four variants of the corresponding optimization problem. We illustrate this with an example.

Example 1: Consider the minimum vertex cover problem, where the input is an L-specification of a graph G . We provide efficient algorithms for the following versions of the problem.

1. **The construction problem:** Output an L-specification of the set of vertices in the approximate vertex cover \mathcal{C} .
2. **The Size problem:** Compute the size of the approximate vertex cover \mathcal{C} for G .
3. **The query problem:** Given any vertex v of G and the path from the root to the node in the *hierarchy tree* (see Section 2 for the definition of hierarchy tree) in which v occurs, determine whether v belongs to the vertex cover \mathcal{C} .
4. **The output problem:** Output the approximate vertex cover \mathcal{C} .

Note that our algorithms for the four variants of the problem apply to the *same* vertex cover \mathcal{C} . Our algorithms for (1), (2) and (3) above run in time **polynomial in the size of the L-specification** rather than the size of the graph obtained by expanding the L-specification.

Our algorithm for (4) runs in time linear in the size of the expanded graph but uses space which is only polynomial in the size of the L-specification. \square

Analogous variants of approximation algorithms can be defined for 1-FPN-specified problems. Therefore, we omit this discussion.

These variants are natural extensions of the definition of approximation algorithms for problems specified using standard specifications. This can be seen as follows: When instances are specified using standard specifications, the number of vertices is polynomial in the size of the description. Given this, any polynomial time algorithm to determine if a vertex v of G is in the approximate minimum vertex cover can be easily modified to obtain a polynomial time algorithm that lists all the vertices of G in the approximate minimum vertex cover. Thus in the case when inputs are specified using standard specifications, (3) can be used to solve (2) and (4) in polynomial time. The above discussion also shows that given an optimization problem specified using standard specifications, variants (1), (3) and (4) discussed above are polynomial time inter-reducible.

The approximation algorithms given in this paper have another desirable feature. For an optimization problem or a query problem, our algorithms use space and time which is a low level polynomial in the size of the hierarchical or the periodic specification. This implies that for graphs of size N , that are specified using specifications of size $O(\text{polylog } N)$, the time and space required to solve problems is only $O(\text{polylog } N)$. Moreover when we need to output the subset of vertices, subset of edges, etc. corresponding to a vertex cover, maximum cut, etc., in the expanded graph, our algorithms take essentially the same time but substantially less (often logarithmically less) space than algorithms that work directly on the expanded graph. The graphs obtained by expanding hierarchical or periodic descriptions are frequently too large to fit into the main memory of a computer [Le86]. This is another reason for designing algorithms which exploit the regular structure of the underlying graphs. Indeed, most of the standard algorithms in the literature assume that the input completely resides in the main memory. As a result, even the most efficient algorithms incur a large number of page faults while executing on the graphs obtained by expanding the hierarchical or periodic specifications. Hence, algorithms designed for solving problems for graphs or circuits represented in a standard fashion are often impractical for succinctly specified graphs. We refer the reader to [Le86, Le90] for more details on this topic.

The rest of the paper is organized as follows. Section 3 contains discussion of related research. In Sections 4, 5 and 6 we give the basic definitions and preliminaries. In Section 7 we discuss our approximation algorithms for L-specified problems and 1-FPN-specified problems. Finally in Section 8, we give concluding remarks and directions for future research.

3 Related research

In the past, much work has been done on characterizing the complexity of various problems when instances are specified using L- or 1-FPN-specifications. For periodically specified graphs, several researchers [CM91, CM93, HW92, KMW67, KO91, Or84a, Or84b] have given

efficient algorithms for solving problems such as determining STRONGLY CONNECTED COMPONENTS, TESTING FOR EXISTENCE OF CYCLES, FINDING MINIMUM COST PATHS BETWEEN A PAIR OF VERTICES, BIPARTITENESS, PLANARITY and MINIMUM COST SPANNING FORESTS. Orlin [Or84b] and Wanke [Wa93] discuss NP- and PSPACE-hardness results for infinite and finite periodically specified graphs.

For L-specified graphs, Lengauer et al. [LW87a, Le88, Le89] and Williams et al. [Wi90] have given efficient algorithms to solve several graph theoretic problems including 2 COLORING, MINIMUM SPANNING FORESTS AND PLANARITY TESTING. Lengauer and Wagner [LW92] show that the following problems are PSPACE-hard when graphs are L-specified: 3 COLORING, HAMILTONIAN CIRCUIT AND PATH, MONOTONE CIRCUIT VALUE PROBLEM, NETWORK FLOW, ALTERNATING GRAPH ACCESSIBILITY and MAXIMUM INDEPENDENT SET. In [LW93], Lengauer and Wanke consider a more general hierarchical specification of graphs based on graph grammars and gave efficient algorithms for several basic graph theoretic problems specified using this specification. We refer the reader to [HLW92, Ma94] for a detailed survey of the work done in the area of hierarchical and periodic specifications.

A substantial amount of research has been done on finding polynomial time approximation algorithms with provable worst case guarantees for NP-hard problems. In contrast, until recently little work has been done towards investigating the existence of polynomial time approximation algorithms for PSPACE-hard problems. As a step in this direction, in [MH+93a, MR+93] we have investigated the existence and non-existence of polynomial time approximations for several PSPACE-hard problems for L-specified graphs. In [HM+94a], we considered geometric intersection graphs defined using the hierarchical specifications (HIL) of Bentley, Ottmann and Widmayer [BOW83]. There, we devised efficient polynomial time approximation schemes for a number of problems for geometric intersection graphs, specified using a restricted form of HIL.

Condon, et al. [CF+93, CF+94] also studied the approximability of several PSPACE-hard optimization problems. They characterize PSPACE in terms of probabilistically checkable debate systems and use this characterization to investigate the existence and non-existence of polynomial time approximation algorithms for a number of basic PSPACE-hard optimization problems.

4 The L-specifications

This section discusses the L-specifications. The following two definitions are essentially from Lengauer [Le89, LW87a, LW92].

Definition 4.1 *An L-specification $\gamma = (G_1, \dots, G_n)$ of a graph is a sequence of labeled undirected simple graphs G_i called cells. The graph G_i has m_i edges and n_i vertices. p_i of the vertices are called pins. The other $(n_i - p_i)$ vertices are called inner vertices. r_i of the inner vertices are called nonterminals. The $(n_i - r_i)$ vertices are called terminals. The remaining $n_i - p_i - r_i$ vertices of G_i that are neither pins nor nonterminals are called explicit vertices.*

Each pin of G_i has a unique label, its name. The pins are assumed to be numbered from 1 to p_i . Each nonterminal in G_i has two labels (v, t) , a name and a type. The type t of a nonterminal in G_i is a symbol from G_1, \dots, G_{i-1} . The neighbors of a nonterminal vertex must be terminals. If a nonterminal vertex v is of the type G_j in G_i , then v has degree p_j and each terminal vertex that is a neighbor of v has a distinct label (v, l) such that $1 \leq l \leq p_j$. We say that the neighbor of v labeled (v, l) matches the l th pin of G_j .

Note that a terminal vertex may be a neighbor of several nonterminal vertices. Given an L-specification γ , $N = \sum_{1 \leq i \leq n} n_i$ denotes the *vertex number*, and $M = \sum_{1 \leq i \leq n} m_i$ denotes the *edge number* of γ . The size of γ , denoted by $\text{size}(\gamma)$, is $N + M$.

Definition 4.2 Let $\gamma = (G_1, \dots, G_n)$ be an L-specification of a graph $E(\gamma)$ and let $\gamma_i = (G_1, \dots, G_i)$. The expanded graph $E(\gamma)$ (i.e. the graph associated with γ) is obtained as follows:

$k = 1$: $E(\gamma) = G_1$.

$k > 1$: Repeat the following step for each nonterminal v of G_k , say of the type G_j : delete v and the edges incident on v . Insert a copy of $E(\gamma_j)$ by identifying the l^{th} pin of $E(\gamma_j)$ with the node in G_k that is labeled (v, l) . The inserted copy of $E(\gamma_j)$ is called a *subcell* of G_k .

Observe that the expanded graph can have multiple edges although none of the G_i have multiple edges. Here however, we **only** consider **simple** graphs, i.e. there is at most edge between a pair of vertices. This means that multi edges are treated simply as single edges. We assume that γ is not redundant in the sense that for each i , $1 \leq i \leq n$, there is a nonterminal v of type G_i in the definition of G_j , $j > i$.

The expansion $E(\gamma)$ is the graph associated with the L-specification γ , with vertex number N . For $1 \leq i \leq n$, $\gamma_i = (G_1, \dots, G_i)$ is the L-specification of the graph $E(\gamma_i)$. Note that the total number of nodes in $E(\gamma)$ can be $2^{\Omega(N)}$. (For example, a complete binary tree with $2^{\Omega(N)}$ nodes can be specified using an L-specification of size $O(N)$.) To each L-specification $\gamma = (G_1, \dots, G_n)$, ($n \geq 1$), we associate a labeled rooted unoriented tree $HT(\gamma)$ depicting the insertions of the copies of the graphs $E(\gamma_j)$ ($1 \leq j \leq n-1$), made during the construction of $E(\gamma)$ as follows:

Definition 4.3 Let $\gamma = (G_1, \dots, G_n)$, ($n \geq 1$) be an L-specification of a graph. The **hierarchy tree** of γ , denoted by $HT(\gamma)$, is the labeled rooted unordered tree defined as follows:

1. Let r be the root of $HT(\gamma)$. The label of r is G_n . The children of r in $HT(\gamma)$ are in one-to-one correspondence with the nonterminal vertices of G_n as follows: The label of the child s of r in $HT(\gamma)$ corresponding to the nonterminal vertex (v, G_j) of G_n is (v, G_j) .
2. For all other vertices s of $HT(\gamma)$ and letting the label of $s = (v, G_j)$, the children of s in $HT(\gamma)$ are in one-to-one correspondence with the nonterminal vertices of G_j as follows: The label of the child t of s in $HT(\gamma)$ corresponding to the nonterminal vertex (w, G_l) of G_j is (w, G_l) .

Given the above definition, we can naturally associate a hierarchy tree corresponding to each γ_i , $1 \leq i \leq n$. We denote this tree by $HT(\gamma_i)$. Note that, each vertex v of $E(\gamma_i)$ is either *an explicit vertex of G_n or is the copy of some explicit vertex v' of G_j ($1 \leq j \leq n$) in exactly one copy C_j^v of the graph $E(\gamma_j)$ inserted during the construction of $E(\gamma_i)$* . This enables us to assign v of $E(\gamma_i)$ to the unique vertex n_v of the $HT(\gamma_i)$ given by

1. if v is a terminal vertex of G_n , then n_v is the root of $HT(\gamma_i)$, and
2. otherwise, v belongs to the node n_v that is the root of the hierarchy tree $HT(\gamma_j)$, corresponding to C_j^v .

Given $HT(\gamma_i)$, the *level number* of a node in $HT(\gamma_i)$ is defined as the length of the path from the node to the root of the tree.

As noted in [Le89], L-specifications have the property that for each copy (instance) of a subcell, a complete boundary description has to be given. Thus if a nonterminal has a lot of pins, copying it is costly. Another property of the definition of L-specifications is that nonterminals are adjacent only to terminals. These properties ensure that the size of the “frontier” (or the number of neighbors) of any nonterminal is polynomial in the size of the specification. These properties weaken the L-specifications with respect to other notions of hierarchy involving a substitution mechanism that entails implicit connections to pins at a cell boundary [Ga82, Wa84]. As a result, regular structures such as grids cannot be specified using small L-specifications. (see [LW87a]). In contrast the graph glueing model of Galperin [Ga82] allows a hierarchical description of pins; thus the size of the frontier can be exponentially large. As a result, graphs such as grids can be represented using descriptions of logarithmic size. However as demonstrated in [Le89, LW87a, Le88, Ga82, Wa84], these properties seem to be a prerequisite for the construction of efficient algorithms for L-specified problems. As subsequent sections show, these restrictions are also necessary in part, for devising efficient approximation algorithms for L-specified problems. The size of the frontier also has a significant impact on the complexity of several basic succinctly specified problems. For example, several basic NP-hard problems become PSPACE-hard when specified using L-specifications (see [LW92, MH+95a]). In contrast, in a recent paper we show that these problems typically become NEXPTIME-hard when specified using the graph glueing specifications of [Ga82] (see [MH+95c]).

By noting Definition 4.1, it follows that an L-specification is a restricted form of context-free graph grammar. The substitution mechanism glues the pins of cells to neighbors of nonterminals representing these cells, as described in Definition 4.2. Such graph grammars are known as *hyperedge replacement systems* [HK87] or *cellular graph grammars* [LW93]. Two additional restriction are imposed on cellular graph grammars to obtain L-specified graphs. First, for each nonterminal there is only one cell that can be substituted. Thus there are no *alternatives* for substitution. Second, the index of the substituted cell has to be *smaller* than the index of the cell in which the nonterminal occurs. The acyclicity condition together with the no alternatives condition implies that an L-specification defines a finite and unique graph.

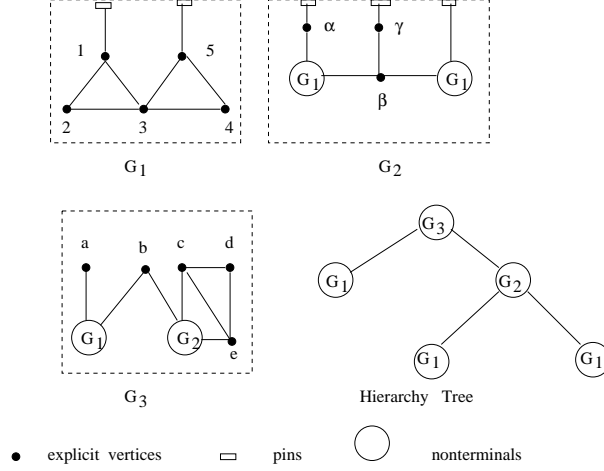


Figure 1: An L-specification G of a graph $E(G)$, and the associated hierarchy tree $HT(G)$. The mapping between the pins and its neighbors is clear by the relative positions of the pins and its neighbors.

We observe that $HT(,)$ is the parse tree of the unique graph generated by the context-free graph grammar $,$.

Example 2: Figures 1 depicts the L-specification $G = (G_1, G_2, G_3)$ and the associate hierarchy tree $HT(G)$. Figure 2 depicts the graph $E(G)$ specified by G . The correspondence between pins of G_j and neighbors of G_j in $G_i, j < i$, is clear by the positions of the vertices and the pins. \square

4.1 Level restricted specifications

Next, we discuss level restricted L-specifications. This is also discussed in [MR+93, MH+93a].

Definition 4.4 An L-specification $, = (G_1, \dots, G_n) (n \geq 1)$, of a graph G is *1-level-restricted*, if for all edges (u, v) of $E(,)$, either

- (1) n_u and n_v are the same vertex of $HT(,)$, or
- (2) one of n_u or n_v is the parent of the other in $HT(,)$.

Extending the above definition we can define k -level-restricted specifications. An L-specification $, = (G_1, \dots, G_n), (n \geq 1)$, of a graph $E(,)$ is *k-level-restricted*, if for all edges (u, v) of $E(,)$, either

- (1) n_u and n_v are the same vertex of $HT(,)$ or
- (2) one of n_u or n_v is an ancestor of the other in $HT(,)$ and the length of the path between n_u and n_v in $HT(,)$ is no more than k .

We note that for any fixed $k \geq 1$, k -level-restricted L-specifications can still lead to graphs that are exponentially large in the sizes of their specifications. Moreover, L-specifications (see




Figure 2: The graph $E(G)$ represented by G specified in Figure 1.

[Le82, Le86, LW87a]) for several practical designs are *k-level restricted* for small values of k . (For example, it is easy to a complete binary tree with $2^{\Omega(N)}$ nodes by a 1-level restricted L-specification of size $O(N)$. Note however, that the specification depicted in Figure 1 is not 1-level-restricted.) For the rest of the paper, given a problem Π specified using standard specifications, we use 1-L- Π to denote the problem specified using 1-level restricted L-specifications and k -L- Π to denote the problem specified using k -level restricted L-specifications.

5 1-FPN-specifications

Next, we give the definition of one dimensional periodic specifications due to Orlin [Or82a], Wanke [Wa93] and Höfting and Wanke [HW92]. For the rest of the paper \mathbf{N} and \mathbf{Z} denotes the set of non-negative integers and integers respectively.

Definition 5.1 *Let $G(V, E)$ (referred to as a static graph) be a finite undirected graph such that each edge (u, v) has an associated non-negative integral weight $t_{u,v}$. The one way infinite graph $G^\infty(V', E')$ is defined as follows:*

1. $V' = \{v(p) \mid v \in V \text{ and } p \in \mathbf{N}\}$
2. $E' = \{(u(p), v(p + t_{u,v})) \mid (u, v) \in E, t_{u,v} \text{ is the weight associated with the edge } (u, v) \text{ and } p \in \mathbf{N}\}$

A 1-dimensional periodic specification, (referred to as 1-PN-specification) is given by $\sigma = (G(V, E))$ and specifies the graph $G^\infty(V', E')$ (referred to as 1-PN-specified graph).

A 1-PN-specification, σ , is said to be narrow or 1-level restricted if $\forall (u, v) \in E, t_{u,v} \in \{0, 1\}$. This implies that $\forall (u(p), v(q)) \in E', |p - q| \leq 1$. Similarly, a 1-PN-specification is k -narrow or k -level restricted if $\forall (u, v) \in E, t_{u,v} \in \{0, 1, \dots, k\}$.




Figure 3: A static graph G , and the graph G^4 specified by the 1-FPN-specification $\gamma = (G, 4)$.

We note that if we replace \mathbf{N} by \mathbf{Z} in Definition 5.1, we obtain a *two way* infinite periodically specified graph defined in Orlin [Or82a]. It is sometimes useful to imagine a narrow periodically specified graph G^∞ as being obtained by placing a copy of the vertex set V at each integral point (also referred to as lattice point) on the X-axis (or the time line) and joining vertices placed on neighboring lattice points in the manner specified by the edges in E .

G^m is the subgraph of the infinite periodic graph G^∞ induced by the vertices associated with nonnegative lattice points less than or equal to m . Formally,

Definition 5.2 Let $G(V, E)$ denote a static graph. Let $G^\infty(V', E')$ denote the one way infinite 1-PN-specified graph as in Definition 5.1. Let $m \geq 0$ be an integer specified using binary numerals. Let $G^m(V^m, E^m)$ be a subgraph of $G^\infty(V', E')$ induced by the vertices $V^m = \{v(p) | v \in V \text{ and } 0 \leq p \leq m\}$. A 1-dimensional finite periodic specification γ (referred to as 1-FPN-specification) is given by $\gamma = (G(V, E), m)$ and specifies the graph G^m (referred to as 1-FPN-specified graph).

An example of a 1-FPN-specified graph appears in Figure 3. In [Or82a], Orlin defined the concept of two way infinite 1-dimensional periodically specified 3CNF formulas and the associated 3SAT problem [GJ79]. It is straightforward to restrict Orlin's definition along the lines of Definition 5.1 to define 1-FPN-specified satisfiability problems. As a consequence, we omit the definition here. (See [Or82a, MH+95a, Pa94] for formal definitions of periodically specified satisfiability problems.) We only give an example of 1-FPN-specified 3CNF formula to illustrate the concept.

Example 3: Let $U = \{x_1, x_2, x_3\}$ be a set of static variables. Let C be a set of static clauses given by $(x_1(0) + x_2(0) + x_3(0)) \wedge (x_1(1) + x_3(0)) \wedge (x_3(1) + x_2(0))$. Let $F = (U, C, 3)$ be a 1-FPN-specification. Then F specifies the 3CNF formula $F^3(U^3, C^3)$ given by

$$\begin{aligned}
& (x_1(0) + x_2(0) + x_3(0)) \wedge (x_1(1) + x_3(0)) \wedge (x_3(1) + x_2(0)) \bigwedge \\
& (x_1(1) + x_2(1) + x_3(1)) \wedge (x_1(2) + x_3(1)) \wedge (x_3(2) + x_2(1)) \bigwedge \\
& (x_1(2) + x_2(2) + x_3(2)) \wedge (x_1(3) + x_3(2)) \wedge (x_3(3) + x_2(2)) \bigwedge \\
& (x_1(3) + x_2(3) + x_3(3))
\end{aligned}$$

6 Other preliminaries

Recall that a graph is said to be *planar* if it can be laid out in the plane in such a way that there are no crossovers of edges. For the rest of the paper, we use L-PL-II, 1-L-PL-II and 1-FPN-PL-II to denote the problem II restricted to L-specified planar instances, 1-level-restricted L-specified planar instances and 1-FPN-specified planar instances respectively. As shown in Lengauer [Le89], given an L-specification Σ , there is a polynomial time algorithm to determine if $E(\Sigma)$ is planar. Similarly as pointed out in [HLW92], given a 1-FPN-specification Σ , there is a polynomial time algorithm to determine if $E(\Sigma)$ is planar. Thus for solving L- or 1-FPN-specified problems restricted to planar instances, we can assume without loss of generality that the inputs to our algorithms consist of planar instances.

Next, we define the problems MAX SAT(**S**). The definition is essentially an extension of the definition of SAT(**S**) given in Schaefer [Sc78].

Definition 6.1 (Schaefer [Sc78])

Let $\mathbf{S} = \{R_1, R_2, \dots, R_m\}$ be a finite set of finite arity Boolean relations. (A Boolean relation is defined to be any subset of $\{0, 1\}^p$ for some integer $p \geq 1$. The integer p is called the **arity** of the relation.) An **S-formula** is a conjunction of clauses each of the form $\hat{R}_i(\xi_1, \xi_2, \dots)$, where ξ_1, ξ_2, \dots are distinct, unnegated variables whose number matches the arity of R_i , $i \in \{1, \dots, m\}$ and \hat{R}_i is the relation symbol representing the relation R_i . The **S-satisfiability** problem is the problem of deciding whether a given **S-formula** is satisfiable.

Given a **S-formula** F , the problem MAX SAT(**S**) is to determine the maximum number simultaneously satisfiable clauses in F .

As in Schaefer [Sc78], given **S**, $\text{Rep}(\mathbf{S})$ is the set of relations that are representable by existentially quantified **S-formulas** with constants.

Recall from [Li82] that a **S-formula** f is said to be planar if its associated bipartite graph is planar. The problem PL-3SAT [Li82] is the problem of determining if a given planar 3CNF formula is satisfiable. Lichtenstein [Li82] showed that the problem PL-3SAT is NP-complete.

Next, we define L-specified **S-formulas**. Such formulas are built by defining larger **S-formulas** in terms of smaller **S-formulas**. Just as L-specifications of graphs can represent graphs that are exponentially larger than the specification, L-specified **S-formulas** can specify formulas that are exponentially larger than the size of the specification.

Definition 6.2 An instance $F = (F_1(X^1), \dots, F_{n-1}(X^{n-1}), F_n(X^n))$ of L-SAT(**S**) is of the form

$$F_i(X^i) = \left(\bigwedge_{1 \leq j \leq l_i} F_{i,j}(X_j^i, Z_j^i) \right) \bigwedge f_i(X^i, Z^i)$$

for $1 \leq i \leq n$ where f_i are **S-formulae**, $X^n = \phi$, X^i, X_j^i, Z^i, Z_j^i , $1 \leq i \leq n-1$, are vectors of Boolean variables such that $X_j^i \subseteq X^i$, $Z_j^i \subseteq Z^i$, $0 \leq i_j < i$. Thus, F_1 is just a **S-formula**. An instance of L-SAT(**S**) specifies a **S-formula** $E(F)$, that is obtained by expanding the F_j , $2 \leq j \leq n$, where the set of variables Z 's introduced in any expansion are considered distinct.

The problem $\text{L-SAT}(\mathbf{S})$ is to decide whether the formula $E(F)$ specified by F is satisfiable. The corresponding optimization problems denoted by $\text{L-MAX-SAT}(\mathbf{S})$ is to find the maximum number of simultaneously satisfiable clauses in $E(F)$.

Let n_i be the total number of variables used in F_i (i.e. $|X^i| + |Z^i|$) and let m_i be the total number of clauses in F_i . The size of F , denoted by $\text{size}(F)$, is equal to $\sum_{1 \leq i \leq n} (m_i n_i)$. Given a formula $E(F)$ specified by an L-specification F , $BG(E(F))$ denotes the bipartite graph associated with $E(F)$. We use $H[BG(E(F))]$ to denote the L-specification of $BG(E(F))$. It is easy to define level-restricted $\text{L-SAT}(\mathbf{S})$ formulas along the lines of Definition 4.4. Hence we omit this definition here.

Example 4: Let $F = (F_1(x_1, x_2), F_2(x_3, x_4), F_3)$ be an instance of L-3SAT where each F_i is defined as follows:

$$F_1(x_1, x_2) = (x_1 + x_2 + z_1) \wedge (z_2 + z_3)$$

$$F_2(x_3, x_4) = F_1(x_3, z_4) \wedge F_1(z_4, z_5) \wedge (z_4 + z_5 + x_4)$$

$$F_3 = F_1(z_7, z_6) \wedge F_2(z_8, z_7)$$

The formula $E(F)$ denoted by F is $(z_7 + z_6 + z_1^1) \wedge (z_2^1 + z_3^1) \wedge (z_8 + z_4 + z_1^2) \wedge (z_2^2 + z_3^2) \wedge (z_4 + z_5 + z_1^3) \wedge (z_2^3 + z_3^3) \wedge (z_4 + z_5 + z_7)$.

We now extend the definition of PL-3SAT given in [Li82] to define the L-PL-3SAT .

Definition 6.3 *The problem L-PL-3SAT is to decide whether the planar 3CNF formula $E(F)$ specified by an L-specification F is satisfiable. The corresponding optimization problems denoted by L-PL-MAX-3SAT is to find the maximum number of simultaneously satisfiable clauses in $E(F)$.*

Extensions of the above definition to 1-L-PL-3SAT , 1-L-PL-MAX-3SAT , $\text{L-PL-SAT}(\mathbf{S})$, $\text{L-PL-MAX-SAT}(\mathbf{S})$, $\text{1-L-PL-SAT}(\mathbf{S})$ and $\text{1-L-PL-MAX-SAT}(\mathbf{S})$ is straightforward and is omitted.

Finally we state the following PSPACE -completeness results proved in a sequel paper [MH+95a]. The definitions of the problems mentioned below can be found in [GJ79].

Theorem 6.4 *The following problems are PSPACE -complete for 1-level-restricted L-specified planar instances: INDEPENDENT SET , DOMINATING SET , VERTEX COVER , $\text{PARTITION INTO TRIANGLES}$ and $\text{SAT}(\mathbf{S})$ such that $\text{Rep}(\mathbf{S})$ is the set of all finite arity Boolean relations.*

Theorem 6.5 *The following problems are PSPACE -complete for 1-FPN-specified planar instances: INDEPENDENT SET , DOMINATING SET , VERTEX COVER , $\text{PARTITION INTO TRIANGLES}$ and $\text{SAT}(\mathbf{S})$ such that $\text{Rep}(\mathbf{S})$ is the set of all finite arity Boolean relations.*

7 Approximation algorithms

The hardness results in Theorems 6.4 and 6.5 motivate the study of polynomial time approximation algorithms with good performance guarantees for these problems. We show that several basic combinatorial problems (including the ones in Theorems 6.4 and 6.5) have approximation algorithms with performance guarantees asymptotically equal to the best known performance guarantees, when instances are specified using standard specifications. As an immediate corollary, most of the problems shown to have polynomial time approximation schemes (PTASs) in [Ba83, HM+94c] when instances are represented using standard specifications, have PTASs when instances are specified either by k -level-restricted L-specifications or 1-FPN-specifications.

7.1 The basic technique: Partial expansion

We outline the basic technique behind the approximation algorithms for the 1-level-restricted L-specified problems. Consider one of the maximization problems Π in this paper. Let \mathcal{A} be an approximation algorithm with performance guarantee $FBEST$, for Π when specified using standard specifications. Also, let $T(N)$ denote an increasing function that is an upper bound on the running time of \mathcal{A} used to solve Π specified using standard specifications of size $O(N)$. Then, given a fixed $l \geq 1$, our approximation algorithm for 1-L- Π takes time $O(N \cdot T(N^{l+1}))$ and has a performance guarantee of $(\frac{l+1}{l}) \cdot FBEST$. Informally, the algorithm consists of $(l + 1)$ iterations. During an iteration i we delete⁶ all the explicit vertices which belong to nonterminals defined at level j , $j = i \bmod (l + 1)$. Intuitively, our algorithm works as follows: break up the given hierarchy tree into a collection of disjoint trees, find a near optimal solution for the vertex induced subgraph⁷ defined by each small tree and output the union of all these solutions as the solution for the problem Π . It is important to observe that the hierarchy tree can have an exponential number of nodes. Hence the deletion of nonterminals and the determination of near optimal solutions for each subtree has to be done in such a manner so that the whole process takes only polynomial time. This is achieved by observing that the subtrees can be divided into n distinct equivalence classes and that the number of subtrees in each equivalence class can be counted in polynomial time in the *size* of the specification.

We remark that our idea of dividing the graph into vertex (edge) disjoint subgraphs is similar to the technique used by Baker [Ba83] for obtaining approximation schemes for planar graph problems.

⁶For minimization problem instead of deleting the vertices in the level, we consider the vertices as a part of both the subtrees.

⁷For a fixed l , the size of each subgraph is polynomial in the size of the specification.

7.2 MAXIMUM INDEPENDENT SET problem for 1-level-restricted L-specified planar graphs

We illustrate the technique by giving a polynomial time approximation scheme for the MAXIMUM INDEPENDENT SET problem for 1-L-specified planar graphs. The INDEPENDENT SET problem is defined as follows. Given a graph $G = (V, E)$ and a positive integer $K \leq |V|$, is there an independent of size K or more for G , i.e., a subset $V' \subseteq V$ with $|V'| \geq K$ such that for each $u, v \in V'$ $(u, v) \notin E$. The optimization problem called the MAXIMUM INDEPENDENT SET PROBLEM (MIS) requires one to find an independent set of maximum size. In [MH+93a], we showed that given an L-specification that has edges between pins in the same nonterminal, there is a polynomial time algorithm to construct a new L-specification such that there is no edge between pins in the *same* nonterminal. Consequently, we assume without loss of generality that in the given L-specification there is no edge between two pins in the same nonterminal.

In the following description, we use $HIS(G_i)$ to denote the approximate independent set for the graph $E(, i)$ obtained by our algorithm **H-MIS**. We also use **F-MIS** to denote the algorithm of Baker [Ba83] for finding an approximate independent set in a planar graph specified using a standard specification. Before we discuss the details of the heuristic we define the concept of *partial expansion* of an L-specification. Recall that, for each nonterminal G_i there is a unique hierarchy tree $HT(G_i)$ rooted at G_i .

Definition 7.1 Let $, = (G_1, \dots, G_n)$ be an L-specification of a graph $E(,)$. The *partial expansion* $PE(G_i^j)$, of the nonterminal G_i is constructed as follows:

$j = 0$: $PE(G_i^j) = G_i - \{\text{all the explicit vertices defined in } G_i\}$ (Thus the definition of $PE(G_i^j)$ now consists of a collection of the nonterminals and pins called in the definition of G_i).

$j \geq 1$: Repeat the following step for each nonterminal G_r called by G_i : Insert a copy of $PE(G_r^{j-1})$ by identifying the l^{th} pin of $PE(G_r^{j-1})$ with the node in G_i that is labeled (v, l) . (Observe that the definition of $PE(G_i^j)$ consists of (i) explicit vertices defined in all the nonterminals at depth r , $0 \leq r \leq j - 1$ in $HT(G_i)$ and (ii) a **multiset** of nonterminals G_k , such that the nonterminal G_k occurs at depth $j + 1$ in the hierarchy tree $HT(G_i)$.)

Let $Ex(PE(G_i^j))$ denote the subgraph induced by the set of explicit vertices in the definition of $PE(G_i^j)$. Also let $V(E(, i))$ denote the set of vertices in $E(, i)$.

Heuristic H-MIS

- **Input:** A 1-level-restricted L-specification $, = (G_1, \dots, G_n)$ of a planar graph G and an integer $l \geq 1$.
- **Output:** An L-specification of an independent set for $E(,)$ whose size is at least $(\frac{l}{l+1})^2$ times the size of an optimal independent set in $E(,)$.

- 1. **For** each $1 \leq i \leq l$, find a near optimal independent set in $E(, i)$ using **F-MIS**.
- 2. **For** each $l + 1 \leq i \leq n - 1$
 - (a) Compute the partial expansion $PE(G_i^l)$ of G_i .
 - (b) Find an independent set in the subgraph $Ex(PE(G_i^l))$ using heuristic **F-MIS**. Denote this by A_i^l .
 - (c) Let G_{i_1}, \dots, G_{i_p} denote the multiset of nonterminals in $PE(G_i^l)$. Then the independent set for the whole graph for the iteration i denoted by $HIS(G_i)$ is given by

$$HIS(G_i) = A_i^l \cup \bigcup_{1 \leq r \leq p} HIS(G_{i_r}).$$

Remark: The explicit vertices in $PE(G_i^l)$ do not have an edge to any of the nonterminals G_{i_1}, \dots, G_{i_p} . From this observation and the definition of hierarchical specification the independent set $HIS(G_i)$ can now be calculated as follows.

(d)

$$|HIS(G_i)| = |A_i^l| + \sum_{1 \leq r \leq p} |HIS(G_{i_r})|$$

- 3. **For** each $0 \leq i \leq l$

- (a) Compute a partial expansion $PE(G_n^i)$ of G_n .
- (b) Find a near optimal independent set of all the explicit vertices in $PE(G_n^i)$ using **F-MIS**. Denote this by A_n^i .
- (c) Let G_{n_1}, \dots, G_{n_p} denote the multiset of nonterminals in $PE(G_n^i)$. The independent set for the whole graph for the iteration i , denoted by $HIS_i(G_n)$, is given by

$$HIS_i(G_n) = A_n^i \cup \bigcup_{1 \leq r \leq p} HIS(G_{n_r}).$$

Remark: By a remark similar to one in Step 2(c) of the algorithm, we have the following.

(d)

$$|HIS_i(G_n)| = |A_n^i| + \sum_{1 \leq r \leq p} |HIS(G_{n_r})|.$$

- 4. The independent set $HIS(G)$ is the largest among all the independent sets $HIS_i(G_n)$ computed in Step 3(c).
- 5. $|HIS(G)| = \max_{0 \leq i \leq l} |HIS_i(G_n)|$

7.3 Analysis and Performance Guarantee

The correctness of **H-MIS** and the proof of its performance guarantee is based on the following intermediate results.

Lemma 7.2 *The set $HIS(G)$ computed by the algorithm **H-MIS** in Step 4 is an independent set.*

Proof: We first prove that the set for $1 \leq i \leq n - 1$, $HIS(G_i)$, is an independent set. The proof is by induction on the depth of the hierarchy tree $HT(\cdot)$.

Basis: If the depth is $\leq l$, the proof follows by the correctness of algorithm **F-MIS**.

Induction: Assume that the lemma holds for all hierarchy trees of depth at most $m > l$. Consider a hierarchy tree of depth $m + 1$. Step 2(c) of the algorithm, computes a partial expansion $PE(G_i^l)$. This implies that the explicit vertices in $PE(G_i^l)$ do not have edges incident on the nonterminals in $PE(G_i^l)$. Thus, by the definition of 1-level-restricted L-specifications and partial expansion, it follows that the independent sets A_i^l , and the sets $HIS(G_{i_r})$, $1 \leq r \leq p$ computed in Steps 2(b) and 2(c) are disjoint. Also, the nonterminals in $PE(G_i^l)$ are at level $l + 1$ in $HT(G_i)$, and have an associated hierarchy tree of depth $\leq m$. Thus by induction hypothesis and the above stated observations, it follows that $HIS(G_m)$ computed in Step 2(c) is an independent set. This completes the proof that $1 \leq i \leq n - 1$, $HIS(G_i)$ is an independent set.

A similar inductive argument proves that the set $HIS_i(G_n)$ computed in each iteration of Step 3(c) is also an independent set. By Step 4, we have that $HIS(G)$ is an independent set.

□

Lemma 7.3 *1. In each iteration i , $l + 1 \leq i \leq n - 1$, of Step 2 of algorithm **H-MIS**, all the explicit vertices in nonterminals at levels $j = l \bmod (l + 1)$ in the hierarchy tree $HT(G_i)$ are deleted.*

*2. In each iteration i of Step 3 of algorithm **H-MIS**, all the explicit vertices in nonterminals at levels $j = i \bmod (l + 1)$ in the hierarchy tree $HT(G_n)$ are deleted.*

Proof of Part 1.: Induction on the depth of the hierarchy tree associated with G_i .

Basis: If the depth is $l + 1$, the proof follows directly by Step 1 and the definition of partial expansion.

Induction: Assume that the lemma holds for all hierarchy trees of depth at most $m > (l + 1)$. Consider a hierarchy tree of depth $m + 1$. Step 2(c) of the algorithm, computes the partial expansion $PE(G_i^l)$. This implies that all the explicit vertices at level l in the hierarchy tree $HT(G_i)$ were deleted. Each nonterminal occurring in the definition of $PE(G_i^l)$ is at level $l + 1$ in $HT(G_i)$, and has an associated hierarchy tree of depth $\leq m$. The proof now follows by induction hypothesis.

Proof of Part 2: Consider a hierarchy tree $HT(G_n)$. In iteration i of Step 3 we compute PE_n^i . This removes all the explicit vertices defined in nonterminals at level i . Also, by the definition of partial expansion it follows that all explicit vertices defined in nonterminals at levels 1 to i appear explicitly in the partially expanded graph. Therefore, the partially expanded graph now has nonterminals defined at level $i + 1$ in the hierarchy tree $HT(G_n)$. The theorem now follows as a consequence of Part 1 of the theorem. □

Given the decomposition of $E(\cdot)$ into a forest (as a result of removing explicit vertices, in nonterminals at levels $j = i \bmod (l + 1)$ in the hierarchy tree $HT(G_n)$) we can associate a hierarchy tree with each of the subgraphs in the forest. Each such tree is a subtree of the

original hierarchy tree $HT(,)$. Label each subtree by the type of nonterminal that is the root of the subtree. The proof of the following lemma is straightforward.

Lemma 7.4 1. During each iteration i of Step 3 of the algorithm **H-MIS**, the root of each subtree is labeled by one of the elements of the set $\{1, \dots, G_{n-1}\}$.

2. For $1 \leq i \leq n$, let $H_1^i, \dots, H_{r_i}^i$ be the set of graphs corresponding to the subtrees labeled G_i . Then for each i the graphs $H_1^i, \dots, H_{r_i}^i$ are isomorphic.

7.3.1 At Least One Good Iteration Exists

Next we prove that, at least one iteration of Step 3 has the property that the number of nodes that are **not** considered in the independent set computation is a small fraction of the optimal independent set.

Let F_i denote the set of vertices obtained by deleting the explicit nodes in iteration i in Step 3 of algorithm **H-MIS**. By Lemma 7.3 it follows that for each iteration i we did not consider the explicit vertices in levels $j_{i_1}, j_{i_2} \dots j_{i_p}$ such that $1 \leq i_p \leq n$ and $j_{i_q} = i \pmod{l+1}$, $1 \leq q \leq p$. Let S_i $0 \leq i \leq l$ be the set of vertices not considered in iteration i of Step 3. Let $IS(G_n)$ denote an optimum independent set in the graph $E(,)$. Let $IS_{opt}(S_i)$ denote the nodes in S_i included in the maximum independent set $IS(G_n)$.

Lemma 7.5

$$\max_{0 \leq i \leq l} |IS(F_i)| \geq \frac{l}{(l+1)} |IS(G_n)|$$

Proof: By Lemma 7.3 and the algorithm **H-MIS**, it follows that

$$S_i \cap S_j = \phi, \quad \cup_{t=0}^{l-1} S_t = V(E(,)), \quad \text{and}$$

$$|IS_{opt}(S_0)| + |IS_{opt}(S_1)| + \dots + |IS_{opt}(S_l)| = |IS(G_n)|.$$

Therefore,

$$\min_{0 \leq i \leq l} |IS_{opt}(S_i)| \leq |IS(G_n)| / (l+1)$$

$$\max_{0 \leq i \leq l} |IS(F_i)| \geq |IS(G_n)| - \min_{0 \leq i \leq l} |IS_{opt}(S_i)| \geq \frac{l}{(l+1)} |IS(G_n)|.$$

□

7.3.2 Performance Guarantee and running time

We now prove that the above algorithm computes a near optimal independent set. Given any $\epsilon > 0$, for some choice of positive integer l such that $(\frac{l}{l+1})^2 \geq (1 - \epsilon)$, we show that algorithm **H-MIS** computes an independent set whose size is at least $(1 - \epsilon)$ times the size of an optimal independent set. We first recall a similar lemma in [Ba83] for planar graphs specified using standard specifications.

Theorem 7.6 [Ba83] *For all fixed $l \geq 1$, given a planar graph G there is linear time algorithm that computes an independent set $FIS(G)$ such that $|FIS(G)| \geq (\frac{l}{l+1}) \cdot |IS(G)|$, where $IS(G)$ denotes a maximum independent set in G .*

Lemma 7.7 $|HIS_i(G_n)| \geq (\frac{l}{l+1}) \cdot |IS(F_i)|$.

Proof: Induction on the number of nonterminals in the definition of $\langle \cdot, \cdot \rangle$. The base case is fairly straightforward. Consider the induction step. By the definition of partial expansion it follows that,

$$|IS(F_i)| = |IS(Ex(PE(G_n^i)))| + \sum_{1 \leq r \leq p} |IS(PE(G_{n_r}))|.$$

From Step 3(c) of the algorithm **H-MIS** we also know that

$$|HIS_i(G_n)| = |A_n^i| + \sum_{1 \leq r \leq p} |HIS(G_{n_r})|.$$

From the induction hypothesis and Theorem 7.6 it follows that

$$\begin{aligned} |A_n^i| &\geq (\frac{l}{l+1}) \cdot |IS(Ex(PE(G_n^i)))| \quad \text{and} \\ |HIS(G_{n_r})| &\geq (\frac{l}{l+1}) \cdot |IS(PE(G_{n_r}))|. \end{aligned}$$

The lemma now follows. \square

Theorem 7.8 $|HIS(G)| \geq (\frac{l}{l+1})^2 \cdot |IS(G)|$.

Proof: Follows from Lemma 7.5 and repeated application of Lemma 7.7. \square

Theorem 7.9 *Let $\langle \cdot, \cdot \rangle$ be an L -specification with vertex number N . Given any $\epsilon > 0$, let $l \geq 1$ be an integer such that $(\frac{l}{l+1})^2 \geq (1 - \epsilon)$. Then the approximation algorithm **H-MIS** runs in time $O(N^{l+2})$ and finds an independent set in $E(\langle \cdot, \cdot \rangle)$ that is at least $(\frac{l}{l+1})^2$ times the size of an optimal independent set in $E(\langle \cdot, \cdot \rangle)$.*

Proof: The performance guarantee follows by Theorem 7.8. Therefore we only prove the claimed time bounds.

First consider Step 1. Note that by Euler's formula, the number of edges in a planar graph with $O(N^l)$ vertices is also $O(N^l)$. Thus, the size of the graphs $E(\langle \cdot, \cdot \rangle_i)$, $1 \leq i \leq l$ is $O(N^l)$. Hence the time required to compute the partial expansion is $O(N^l)$. By Theorem 7.6, the

time needed to compute an independent set in $E(, i)$ is $O(N^l)$. Thus the total running time of Step 1 is $O(N^l)$.

Next consider each iteration of Step 2 of the algorithm **H-MIS**. Step 2(a) takes time $O(N^{l+1})$ since the size of the graph $PE(G_i^l)$ can be $O(N^{l+1})$. By Theorem 7.6, the time needed for executing Step 2(b) is $O(N^l)$, since the number of nodes in $Ex(PE(G_i^l))$ can be $O(N^l)$. By Lemma 7.4, Step 2(c) and 2(d) together take time $O(N)$. Thus the total running time for executing one iteration of Step 2 is $O(N^{l+1})$. Thus the total running time of Step 2 is $nO(N^{l+1}) = O(N^{l+2})$.

A similar calculation shows that the total time needed to execute one iteration of Step 3 is $O(N^{l+1})$. Thus the total time needed to execute Step 3 is $(l+1)O(N^{l+1}) = O(N^{l+1})$.

Thus the total running time of the algorithm is $O(N^{l+2})$. \square

7.4 L-Specification of the solution and the query problem

In Section 7.3, we showed how to solve the **size** problem for 1-L-MIS. We now discuss the **construction** problem. As noted in Section 2.1 our algorithms for the four variants of the problem apply to the *same* independent set $HIS(G)$.

The L-specification of the solution can be easily constructed by slightly modifying the algorithm **H-MIS** as follows. Consider the iteration i of Step 3 which gives the maximum independent set. Denote the iteration by i^* . The L-specification H of the solution consists of nonterminals H_1, \dots, H_n . For $1 \leq j \leq n$ the explicit vertices of H_j are the explicit vertices in $PE(G_j^i)$ that are in the independent set. If $PE(G_j^i)$ calls nonterminals G_{j_1}, \dots, G_{j_m} then the nonterminal H_j calls the nonterminals H_{j_1}, \dots, H_{j_m} . Observe that some of the nonterminals H_i may be redundant and these can be removed from the final specification. Given the L-specification of the solution, the **query** problem can be easily solved by examining if the given vertex occurs in the set of nodes specified by the L-specification of the solution. Given an L-specification of the solution, we can solve the **output** problem as follows: We traverse the hierarchy tree associated with H in a depth first manner and output the vertices in the nonterminals visited during the traversal.

Observe that the only place we used planarity was to obtain a near optimal solution for the maximum independent set problem for each partially expanded graph. In Section 7.7 we use this observation to compute near optimal solutions for problems for arbitrary 1-level-restricted L-specified graphs.

7.5 Other L-specified planar problems

Our technique can be applied to obtain efficient approximation algorithms for the following additional optimization problems: MINIMUM VERTEX COVER, MINIMUM DOMINATING SET, MAXIMUM PARTITION INTO TRIANGLES, MINIMUM EDGE DOMINATING SET, MAXIMUM CUT and MAX SAT(**S**) for any finite set of finite set of finite arity Boolean relations **S**. The basic idea behind devising approximation schemes for these problems is similar to the ideas used

to solve the MAXIMUM INDEPENDENT SET problem. Therefore, we only briefly discuss the method for MINIMUM VERTEX COVER and MAX SAT(**S**).

(1) MINIMUM VERTEX COVER:

Given a graph $G = (V, E)$ and a positive integer $K \leq |V|$, is there a vertex cover of size K or less for G , i.e., a subset $V' \subseteq V$ with $|V'| \leq K$ such that for each edge $(u, v) \in E$ either u or v belongs to V' . The optimization problem requires one to find a vertex cover of minimum size.

In order to approximate the 1-L-PL-MINIMUM VERTEX COVER problem we do the following. Given an ϵ , we choose an l such that $(\frac{l+1}{l})^2 \leq (1+\epsilon)$. Next, we modify the definition of partial expansion so that instead of deleting the explicit vertices at levels $(l+1)$ apart, we consider them in both sides of the partition. For each $0 \leq i < l$, the algorithm finds a near optimal solution for the overlapping planar graphs induced by explicit vertices in levels $(jl+i)$ to $((j+1)l+i)$, for $j \geq 0$. The algorithm picks the best among all the vertex covers obtained for the different values of i . Let $OPT(G)$ denote an optimal vertex cover for G . The following lemma points out that the solution obtained is at most $(\frac{l+1}{l})^2$ times the optimal vertex cover. The proof of the lemma follows the same general argument given for the MAXIMUM INDEPENDENT SET problem.

Lemma 7.10 *The size of the vertex cover obtained is no more than $(\frac{l+1}{l})^2 |OPT(G)|$*

Proof: Consider an optimal solution $OPT(G)$ to the vertex cover problem. Then for some $0 \leq t < l$, at most $|OPT(G)|/l$ nodes in $OPT(G)$ are in levels congruent to $t \bmod (l)$. Consider the iteration when the planar graphs are obtained by overlapping at levels congruent to $t \bmod (l)$. Hence the size of an optimal vertex cover in this iteration is $(|OPT(G)| + |OPT(G)|/l)$. Now applying the known approximation scheme [Ba83] for computing a near optimal vertex cover for each smaller subgraphs, we obtain a near optimal vertex cover for the whole graph for iteration t . The size of the vertex cover obtained in this iteration is no more than $(|OPT(G)| + |OPT(G)|/l) \frac{l+1}{l}$. The reason for this is that the explicit vertices in the overlapping levels are counted twice and the near optimal vertex cover heuristic for the each of the small graphs yield a vertex cover of size $(l+1)/l$ times the optimal vertex cover for each small graph. Since the heuristic picks the minimum vertex over all values of i , it follows that the size of the vertex cover produced by the heuristic is no more than $(\frac{l+1}{l})^2 |OPT(G)|$. \square

(2) MAX SAT(S**):**

In the following, we will assume that an instance F of 1-L-PL-MAX-SAT(**S**) is specified by $H[BG(E(F))]$ (i.e the specification of the associated bipartite graph.). The basic idea behind the approximation schemes for 1-L-MAX-PL-SAT(**S**) is as follows: For each i , $0 \leq i \leq 2l$ in increments of 2, we remove the explicitly defined clauses which are in levels j and $j+1$, such that $j = i \bmod (l+1)$. This breaks the bipartite graph into a number of smaller bipartite graphs such that the formulas they denote do not share any variables or clauses. It is not

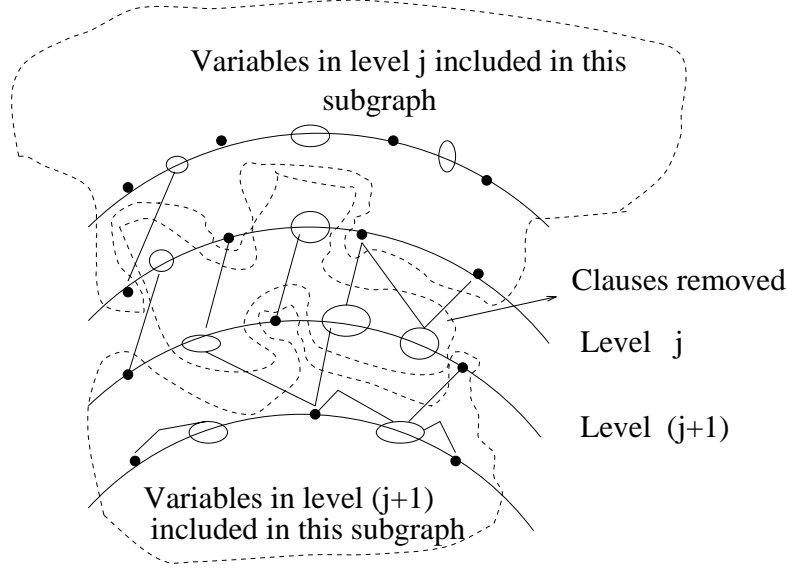


Figure 4: Basic Idea behind the approximation algorithm for 1-L-MAX-PL-SAT(**S**). The black dots represent variables and the ellipses denote clauses. The figure depicts which set of clause to remove and the redistribution of the variables.

difficult to modify the definition of partial expansion to obtain a decomposition as described above. Figure 4 shows how the variables in levels j and $j + 1$ are redistributed. As in the case of MAXIMUM INDEPENDENT SET problem, it is easy to see that there exists an iteration t , $0 \leq t \leq 2l$, such that at most $\frac{OPT}{(l+1)}$ clauses in OPT are deleted. Next, by the results in [HM+94c] the problem can be solved near optimally for each smaller subformulas. The union of the clauses satisfied for each small formula constitutes a solution for a given value of i . We pick the best solution for different values of i . This ensures that the best assignment to the variables over all values of i is at least $(\frac{l}{l+1})^2$ of an optimal assignment to the variables of the 1-L-PL-MAX-SAT(**S**) instance.

7.6 Extension to k -level restricted instances

The above results can be generalized for problems specified using k -level-restricted L-specifications. We only point out the essential differences. Again, for the purposes of illustration consider the problem k -L-PL-MIS. First note that we need to extend the definition of partial expansion so that we delete the explicit vertices in nonterminals at k consecutive levels. This implies that the time to compute $PE(G_i^l)$, $1 \leq i \leq n - 1$ is $O(N^{l+k})$. The rest of the algorithm follows the same outline as that of **H-MIS**. The proof of correctness and the performance guarantee also follow similar arguments as in Section 7.3. Thus the total running time of the algorithm is $O(N^{k+l+1})$ and its performance guarantee is $(\frac{l+1}{l})^2$. Hence we have the following theorem.

Theorem 7.11 *For any fixed $k \geq 1$, there are polynomial time approximation schemes for the problems MAXIMUM INDEPENDENT SET, MINIMUM DOMINATING SET, MINIMUM VERTEX*

COVER, MINIMUM EDGE DOMINATING SET, MAXIMUM PARTITION INTO TRIANGLES AND MAXIMUM CUT, and MAX SAT(**S**), for each finite set of finite arity Boolean relations **S**, when restricted to planar instances specified using k -level-restricted L -specifications.

7.7 Extension to arbitrary graphs

Our results in Sections 7.2 through 7.6 can be extended for problems on arbitrary graphs specified using k -level-restricted L -specifications. To do this, observe that to obtain the results in Sections 7.2 through 7.6 we used planarity only to obtain approximation schemes for smaller subgraphs (formulas) obtained as a result of partial expansion. If the graphs were not planar we could use the best known approximation algorithms for solving the problem near optimally and in turn get a performance guarantee which reflects this bound. For example, consider the problem 1-L-MAX-2SAT. Let $\epsilon > 0$ be the required performance guarantee. $l \geq 1$ is an integer satisfying the inequality $\frac{l}{l+1} \geq (1 - \epsilon)$. For the problem MAX-2SAT, the recent work of Goemans and Williamson [GW94] provides an approximation algorithm with performance guarantee of 1.137. Using their algorithm as a subroutine to solve the small MAX-2SAT instances obtained as a result of partial expansion, we can devise an approximation algorithm for 1-L-MAX-2SAT with performance guarantee $\left(\frac{l+1}{l}\right) 1.137$. A similar idea applies to other optimization problems considered. Again, it is easy to generalize our results for k -level restricted L -specifications. Thus we have the following theorem.

Let Π be one of the problems: MAXIMUM INDEPENDENT SET, MINIMUM DOMINATING SET, MINIMUM VERTEX COVER, MINIMUM EDGE DOMINATING SET, MAXIMUM PARTITION INTO TRIANGLES, MAXIMUM CUT and MAX-SAT(**S**), for finite set of Boolean relations **S**, such that $Rep(\mathbf{S})$ is the set of all finite arity Boolean relations ⁸.

Theorem 7.12 *For all fixed $k \geq 1$, $\epsilon > 0$ and for all of the problems Π , there are polynomial time approximation algorithms with performance guarantee⁹ $(1 + \epsilon) \cdot FBEST_{\Pi}$ for problems Π , when specified using k -level restricted L -specifications. Here $FBEST_{\Pi}$ denotes the best known performance guarantee of an algorithm for the problem Π for instances specified using standard specifications.*

Using the results of Arora et al.[AL+92], Bellare et. al. [BG+95] and our results in [HMS94] we get the following theorem.

Theorem 7.13 *Unless $P = NP$, the problems Π , when specified using k -level restricted L -specifications, do not have polynomial time approximation schemes.*

7.8 Approximation algorithms for 1-FPN-specified problems

Next, we briefly discuss how to extend our ideas developed in Sections 7.2 through 7.7 in order to devise approximation schemes for several PSPACE-hard problems for 1-FPN-specified instances.

⁸Actually our easiness results hold for all finite set of finite arity Boolean relations **S**.

⁹For the sake of uniformity we assume that the performance guarantee is ≥ 1 .




Figure 5: A schematic diagram showing the vertices to be removed in each iteration i while computing a near optimal independent set for 1-FPN-specified planar graphs. In our example $i = 3$, $l + 1 = 4$, and $m = 9$. Each box represents a copy of the vertices in the original static graph. The shaded area represents the vertices that are removed.

The basic idea is simple. Once again we illustrate our ideas by describing our approximation algorithm for the problem 1-FPN-PL-MIS. Given a 1-FPN-specification $\gamma = (G(V, E), m)$ of a planar graph G^m and an $\epsilon > 0$, we find the corresponding integer l that satisfies the inequality $(\frac{l}{l+1})^2 \geq (1 - \epsilon)$. For $0 \leq i \leq l$, we remove the vertices placed at the lattice points j such that $j = i \bmod(l + 1)$. This partitions the graph G^m into a number of smaller disjoint subgraphs, each induced by l consecutive lattice points.

Specifically, for a given i , let $l_p^i = \max\{0, (p - 1)(l + 1) + (i + 1)\}$ and $r_p^i = \min\{m, p(l + 1) + (i - 1)\}$, where $0 \leq p \leq t_i$. Here $t_i = \lceil \frac{m - (i - 1)}{l + 1} \rceil$. Let the subgraph induced by vertices $v(j_p)$, where $l_p^i \leq j_p \leq r_p^i$, be denoted by $H(l_p^i, r_p^i)$. For a given $\epsilon > 0$, the graphs $H(l_p^i, r_p^i)$ are linear in the size of γ . Figure 5 shows a schematic diagram of the vertices removed in a given iteration i . Next, we solve the MIS problem near optimally on each of the subgraphs. This can be done by using the linear time algorithm stated in Theorem 7.6. The union of these independent sets is the independent set obtained in iteration i . The heuristic simply picks up the largest independent set obtained over all $l + 1$ iterations. By arguments similar to the ones we presented for approximating 1-L-PL-MIS (Subsections 7.2 to 7.4), it follows that the approximation algorithm has a performance guarantee of $(\frac{l+1}{l})^2$.

We note the following important point. If a near optimal independent set were to be obtained for each subgraph $H(l_p^i, r_p^i)$, we would take an exponential amount of time in each iteration i . This is because $p = O(m)$. Hence we can not afford to solve the problem explicitly for each subgraph. But observe that each iteration i the subgraphs $H(l_p^i, r_p^i)$, $1 \leq p \leq \lceil \frac{m - (i - 1)}{l + 1} \rceil - 1$ are isomorphic. Hence we need to solve the MIS problem for the graphs $H(l_0^i, r_0^i)$, $H(l_1^i, r_1^i)$ and $H(l_{t_i}^i, r_{t_i}^i)$, where $t_i = \lceil \frac{m - (i - 1)}{l + 1} \rceil$. Let $IS(H(l_p^i, r_p^i))$ denote the independent set obtained by the heuristic for the graph $H(l_p^i, r_p^i)$. Furthermore, let the approximate maximum independent set for the whole graph for a given iteration i be denoted by $IS(G^m(i))$. Then $IS(G^m(i))$ is given by the following equation:

$$|IS(G^m(i))| = |IS(H(l_0^i, r_0^i))| + \lfloor \frac{m - (i - 1)}{l + 1} \rfloor |IS(H(l_1^i, r_1^i))| + |IS(H(l_{t_i}^i, r_{t_i}^i))|$$

This completes the discussion of the approximation algorithm for 1-FPN-PL-MIS. By combining the above arguments along with those in Sections 7.2 through 7.7, we can show that several other optimization problems can be approximated in a similar fashion. Again it is easy

to see that the technique extends to problems for arbitrary instances and also to problems for instances specified using k -narrow 1-FPN-specifications. Thus we have the following theorem.

Theorem 7.14 *For all fixed $k \geq 1$, $\epsilon > 0$ and for all of the problems Π stated in Section 7.7, there are polynomial time approximation algorithms with performance guarantee¹⁰ $(1 + \epsilon) \cdot FBEST_{\Pi}$ for problems Π , when specified using k -level restricted 1-FPN-specifications. Here $FBEST_{\Pi}$ denotes the best known performance guarantee of an algorithm for the problem Π for instances specified using standard specifications.*

Observe that the technique used to devise approximation algorithms for problems restricted to k -narrow 1-FPN-specified instances is very similar to the technique used to devise approximation algorithms for k -level-restricted L-specified problems. But there are two important differences in the details of the algorithms.

1. In case of algorithms for L-specified problems, the number of equivalence classes is $O(n)$ where n is the number of nonterminals. In contrast, the number of equivalence classes in case of algorithms for 1-FPN-specified problems is only $O(1)$.
2. The size of the subgraphs for which the problem is solved near optimally also differs significantly. Specifically, the number of explicit vertices in $PE(G_i^l)$ can be $O(N^l)$. Moreover the time required to compute $PE(G_i^l)$ can be $O(N^{l+k})$. In contrast, the number of explicit vertices in each $H(l_p^i, r_p^i)$ is only $O(N)$ and the time required to construct each $H(l_p^i, r_p^i)$ is only $O(N)$. In both cases we use N to be the vertex number of the respective specifications, $(,)$ (N can be $O(size(,))$).

These important differences allow us to devise linear time approximation schemes for 1-FPN-specified problems.

8 Conclusions

8.1 Summary

We have investigated the polynomial time approximability of several PSPACE-hard optimization problems for both L- and 1-FPN-specified instances. A general approach was given to obtain polynomial time approximation schemes for several PSPACE-hard optimization problems for planar graphs specified using k -level restricted L- or 1-FPN-specifications.

In [HM+94a], we used the technique of partial expansion to obtain polynomial time approximation schemes for geometric intersection graphs specified using a restricted form of the hierarchical specifications proposed by Bentley, Ottmann and Widmayer [BOW83]. We believe that the partial expansion technique can be used to obtain efficient approximations for other problems specified using L- or 1-FPN-specifications as well as for problems specified using other succinct specifications.

¹⁰For the sake of uniformity we assume that the performance guarantee is ≥ 1 .

In an accompanying paper [MH+95a], we investigate the decision complexity of various combinatorial problems specified using various kinds of L-specifications and 1-FPN-specifications. There we give a general method to obtain PSPACE-hard lower bounds for such problems including the ones discussed here.

8.2 Open Problems

We conclude with a list of open problems for future research.

1. Can we use the concept of Debate systems [CF+93, CF+94] to prove non-approximability results for problems specified using arbitrary (not level-restricted) L-specifications ?

Recently, Agarwal and Condon [AC95] have partially answered this question by showing that unless $P = PSPACE$, there is no polynomial time approximation scheme for the problem L-MAX-3SAT. The result was proved by using the characterization of PSPACE in terms of random debate systems. In [HMS94], we extended their result to hold for any L-MAX-SAT(**S**) such that $Rep(\mathbf{S})$ denotes the set of all finite arity Boolelan relations.

2. Recently, several researchers have considered logical definability of a number of optimization problems and defined appropriate classes such as MAX SNP MAX Π_1 MAX NP and MAX #P (cf. [KT94, Kan92b, PY91, PR93]). All these researchers have assumed the the input is specified using standard specifications. What happens if the instances (finite or infinite) are specified succinctly ?

Some work has been done along these lines by Hirst and Harel [HH93]. Specifically, they considered infinite recursive versions of several NP optimization problems. They prove that some problems become highly undecidable (in terms of Turing degrees) while others remain on low levels of arithmetic hierarchy. As a corollary of their results they provide a method for proving (finitary) problems to be outside the syntactic class MAX NP and hence outside MAX SNP.

Acknowledgements: We thank the referees for their invaluable comments. We also thank Anne Condon, Ashish Naik, Egon Wanke, Joan Feigenbaum, R. Ravi, S.S. Ravi and Thomas Lengauer for many helpful conversations during the course of writing this paper.

References

- [AC95] S. Agarwal and A. Condon, “On Approximation Algorithms for Hierarchical MAX-SAT,” *Proc. 10th IEEE Annual Conference on Structure in Complexity Theory*, June 1995, pp. 181-190.
- [AL+92] S. Arora, C. Lund, R. Motwani, M. Sudan and M Szegedy, “Proof Verification and Hardness of Approximation Problems,” *Proc. 33rd IEEE Symposium on Foundations of Computer Science (FOCS)*, 1992, pp. 14-23.
- [Ba83] B.S. Baker, “Approximation Algorithms for NP-Complete Problems on Planar Graphs,” *Journal of the ACM (J. ACM)*, Vol. 41, No. 1, 1994, pp. 153-180.
- [BG+95] M. Bellare, O. Goldreich and M. Sudan, “Free Bits, PCPs and Non-Approximability – Towards Tight Results,” *Proc. 36rd IEEE Symposium on Foundations of Computer Science (FOCS’95)*, Oct. 1995, pp. 422-431.
- [BOW83] J.L. Bentley, T. Ottmann and P. Widmayer, “The Complexity of Manipulating Hierarchically Defined set of Rectangles,” *Advances in Computing Research*, ed. F.P. Preparata, Vol. 1, (1983), pp. 127-158.
- [CM91] E. Cohen and N. Megiddo, “Recognizing Properties of Periodic graphs,” *Applied Geometry and Discrete Mathematics*, Vol. 4, *The Victor Klee Festschrift*, P. Gritzmann and B. Strumfels, eds., ACM, New York, 1991, pp. 135-146.
- [CM93] E. Cohen and N. Megiddo, “Strongly Polynomial-time and NC Algorithms for Detecting Cycles in Dynamic Graphs,” *Journal of the ACM (J. ACM)* Vol. 40, No. 4, September 1993, pp. 791-830.
- [CF+93] A. Condon, J. Feigenbaum, C. Lund and P. Shor, “Probabilistically Checkable Debate Systems and Approximation Algorithms for PSPACE-Hard Functions,” in *Proc. 25th ACM Symposium on Theory of Computing (STOC)*, 1993, pp. 305-313.
- [CF+94] A. Condon, J. Feigenbaum, C. Lund and P. Shor, “Random Debators and the Hardness of Approximating Stochastic functions,” *Proc. 9th IEEE Annual Conference on Structure in Complexity Theory*, June 1994, pp. 280-293.
- [Ga59] D. Gale, “Transient Flows in Networks,” *Michigan Mathematical Journal*, No. 6, 1959 , pp. 59-63.
- [Ga82] H. Galperin “Succinct Representation of Graphs,” Ph.D. Thesis, Princeton University, 1982.
- [GJ79] M.R. Garey and D.S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-Completeness*, Freeman, San Francisco CA, 1979.
- [GJM91] C. Ghezzi, M. Jazayeri and D. Mandrioli, *Fundamentals of Software Engineering*, Prentice Hall, Englewood Cliffs, NJ, 1991.
- [GW94] M.X. Goemans and D.P. Williamson “878 Approximation Algorithms for MAX CUT and MAX 2SAT,” *Proc. 26th Annual ACM Symposium on Theory of Computing, (STOC)*, May 1994, pp. 422-431. To appear in *Journal of the ACM (J. ACM)*.

- [HK87] A. Habel and H.J. Kreowski, "May we Introduce to you: Hypergraph Languages Generated by Hyperedge Replacement," *Proc. 13th International Workshop on Graph-Theoretic Concepts in Computer Science (WG'87)*, Springer Verlag, LNCS, Vol. 291, 1987, pp. 15-26.
- [Ha75] F.O. Hadlock, "Finding a Maximum Cut in a Planar Graph in Polynomial Time," *SIAM Journal on Computing*, No. 4, 1975, pp. 221-225.
- [HH93] T. Hirst, D. Harel, "Taking it to the Limit: On Infinite Variants of NP-Complete Problems," *Proc. 8th IEEE Annual Conference on Structure in Complexity Theory (Structures' 93)*, June, 1993, pp. 292-304.
- [HLW92] F. Höfting, T. Lengauer and E. Wanke, "Processing of Hierarchically Defined Graphs and Graph Families," *Data Structures and Efficient Algorithms* (Final Report on the DFG Special Joint Initiative), Springer-Verlag, LNCS 594, 1992, pp. 44-69.
- [HW92] F. Höfting and E. Wanke, "Minimum Cost Paths in Periodic Graphs," *SIAM Journal on Computing*, Vol. 24, No. 5, 1995, pp. 1051-1067.
- [HM+94a] H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, S. S. Ravi, D. J. Rosenkrantz and R. E. Stearns, "A Unified Approach to Approximation Schemes for NP- and PSPACE-Hard Problems for Geometric Graphs," *Proc. 2nd Annual European Symposium on Algorithms (ESA'94)*, September, 1994, pp. 424-435.
- [HM+94c] H.B. Hunt III, M.V. Marathe, V. Radhakrishnan, D.J. Rosenkrantz and R.E. Stearns, "Designing Approximation Schemes Using L-Reductions," in *Proc. of the 14th Annual Foundations of Software Technology and Theoretical Computer Science (FST & TCS)*, Madras, India, December, 1994, pp. 342-353.
- [HMS94] H.B. Hunt III, M.V. Marathe and R.E. Stearns, "Generalized CNF Satisfiability Problems and Non-Efficient Approximability," *Proc. 9th IEEE Conf. on Structure in Complexity Theory*, June-July 1994, pp. 356-366. A detailed version of the paper as University at Albany Technical Report, TR-95-27, May 1995.
- [IS87] K. Iwano and K. Steiglitz, "Testing for Cycles in Infinite Graphs with Periodic Structure," *Proc. 19th Annual ACM Symposium on Theory of Computing, (STOC)*, 1987, pp. 46-53.
- [IS88] K. Iwano and K. Steiglitz, "Planarity Testing of Doubly Connected Periodic Infinite Graphs," *Networks*, No. 18, 1988, pp. 205-222.
- [Kan92b] V. Kann, "On the Approximability of NP-complete Optimization Problems," Ph.D. Thesis, Dept. of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm, Sweden, May 1992.
- [KMW67] R.M. Karp, R.E. Miller and S. Winograd, "The Organization of Computations for Uniform Recurrence Equations," *Journal of the ACM (J. ACM)*, Vol. 14, No. 3, 1967, pp. 563-590.
- [KO91] M. Kodialam and J.B. Orlin, "Recognizing Strong Connectivity in Periodic graphs and its relation to integer programming," *Proc. 2nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1991, pp. 131-135.
- [KT94] P. G. Kolaitis and M.N. Thakur, "Logical Definability of NP Optimization Problems," *Information and Computation*, No. 115, 1994, pp. 321-353.

- [KS88] K. R. Kosaraju and G.F. Sullivan, "Detecting Cycles in Dynamic Graphs in Polynomial Time," *Proc. 29th IEEE Symposium on Foundations of Computer Science (FOCS)*, 1988, pp. 398-406.
- [Le82] T. Lengauer, "The Complexity of Compacting Hierarchically Specified Layouts of Integrated Circuits," *Proc. 23rd IEEE Symposium on Foundations of Computer Science (FOCS)*, 1982, pp. 358-368.
- [Le86] T. Lengauer, "Exploiting Hierarchy in VLSI Design," *Proc. AWOC '86*, Springer Verlag, LNCS 227, 1986, pp. 180-193.
- [LW87a] T. Lengauer and E. Wanke, "Efficient Solutions for Connectivity Problems for Hierarchically Defined Graphs," *SIAM Journal on Computing*, Vol. 17, No. 6, 1988, pp. 1063-1080.
- [LW87b] T. Lengauer and C. Weiner, "Efficient Solutions Hierarchical Systems of Linear Equations," *Computing*, Vol 39, 1987, pp. 111-132.
- [Le88] T. Lengauer, "Efficient Algorithms for Finding Minimum Spanning Forests of Hierarchically Defined graphs," *Journal of Algorithms*, Vol. 8, 1987, pp. 260-284.
- [Le89] T. Lengauer, "Hierarchical Planarity Testing," *Journal of the ACM (J. ACM)*, Vol. 36, No. 3, July 1989, pp. 474-509.
- [Le90] T. Lengauer, *Combinatorial Algorithms for Integrated Circuit Layout*, John Wiley and Sons, 1990.
- [LW92] T. Lengauer and K.W. Wagner, "The Correlation Between the Complexities of Non-Hierarchical and Hierarchical Versions of Graph Problems," *Journal of Computer and System Sciences (JCSS)*, Vol. 44, 1992, pp. 63-93.
- [LW93] T. Lengauer and E. Wanke, "Efficient Decision Procedures for Graph Properties on Context-Free Graph Languages," *Journal of the ACM (J. ACM)*, Vol. 40, No. 2, 1993, pp. 368-393.
- [Li82] D. Lichtenstein, "Planar Formulae and their Uses," *SIAM Journal on Computing*, Vol 11, No. 2, May 1982 , pp. 329-343.
- [MH+93a] M.V. Marathe H.B. Hunt III, and S.S. Ravi, "The Complexity of Approximating PSPACE-Complete Problems for Hierarchical Specifications," *Nordic Journal of Computing*, Vol. 1, 1994, pp. 275-316.
- [MR+93] M.V. Marathe, V. Radhakrishnan, H.B. Hunt III, and S.S. Ravi, "Hierarchical Specified Unit Disk Graphs," *Proc. 19th International Workshop on Graph-Theoretic Concepts in Computer Science (WG '93)*, June, 1993, pp. 21-32. To appear in *Theoretical Computer Science*.
- [Ma94] M.V. Marathe, *Complexity and Approximability of NP- and PSPACE-hard Optimization Problems*, Ph.D. thesis, Department of Computer Science, University at Albany, Albany, NY August, 1994.
- [MH+95a] M.V. Marathe, H.B. Hunt III, R.E. Stearns and V. Radhakrishnan, "Complexity of Hierarchically and 1-Dimensional Periodically Specified Problems," Technical Report LAUR-93-3348, Los Alamos National Laboratory, August, 1995. Also to be presented at the *DIMACS Workshop on Satisfiability Problem: Theory and Applications*, March 1996.

- [MH+95c] M.V. Marathe, H.B. Hunt III and R.E. Stearns “A Uniform Approach to prove NEXPTIME-hardness for problems specified by G.C.R, S.C.R and BOW Specifications,” manuscript, November 1995.
- [MTM92] J.O. McClain, L.J. Thomas and J.B. Mazzola, *Operations Management*, Prentice Hall, Englewood Cliffs, 1992.
- [MC80] C. Mead and L. Conway, *Introduction to VLSI systems*, Addison Wesley, 1980.
- [Or82a] J.B. Orlin, “The Complexity of Dynamic/Periodic Languages and Optimization Problems,” Sloan W.P. No. 1679-86, July 1985, Working paper, Alfred P. Sloan School of Management, MIT, Cambridge, MA 02139. A Preliminary version of the paper appears in *Proc. 13th ACM Annual Symposium on Theory of Computing (STOC)*, 1981, pp. 218-227.
- [Or84a] J.B. Orlin, “Maximum Convex Cost Dynamic Network Flows,” *Mathematics for Operations Research*, Vol. 9, No. 2, May 1984, pp. 190-206.
- [Or84b] J.B. Orlin, “Some Problems on Dynamic/Periodic Graphs,” *Progress in Combinatorial Optimization*, Academic Press, May 1984, pp. 273-293.
- [PR93] A. Panconesi and D. Ranjan, “Quantifiers and Approximations,” *Theoretical Computer Science*, 107, 1993, pp. 145-163.
- [PY86] C. Papadimitriou and M. Yannakakis, “A note on Succinct Representation of Graphs,” *Information and Computation*, No. 71, 1986, pp. 181-185.
- [PY91] C. Papadimitriou and M. Yannakakis, “Optimization, Approximation and Complexity Classes,” *Journal of Computer and System Sciences (JCSS)*, No. 43, 1991, pp. 425-440.
- [Pa94] C. Papadimitriou, *Computational Complexity*, Addison-Wesley, Reading, Massachusetts, 1994.
- [RH93] D.J. Rosenkrantz and H.B. Hunt III, “The Complexity of Processing Hierarchical Specifications,” *SIAM Journal on Computing*, Vol. 22, No. 3, 1993, pp. 627-649.
- [Sc78] T. Schaefer, “The Complexity of Satisfiability Problems,” *Proc. 10th ACM Symposium on Theory of Computing (STOC)*, 1978, pp. 216-226.
- [Wa84] K.W. Wagner, “The Complexity of Problems Concerning Graphs with Regularities,” *Proc. 11th Symposium on Math. Foundations of Computer Science (MFCS)*, LNCS 176, Springer-Verlag, 1984, pp. 544-552.
- [Wa93] E. Wanke, “Paths and Cycles in Finite Periodic Graphs,” *Proc. 20th Symposium on Math. Foundations of Computer Science (MFCS)*, LNCS 711, Springer-Verlag, 1993, pp. 751-760.
- [Wi90] M. Williams, “Efficient Processing of Hierarchical Graphs,” *TR 90-06*, Dept of Computer Science, Iowa State University. (Parts of the report appeared in WADS’89, pp. 563-576 and SWAT’90, pp. 320-331 coauthored with Fernandez-Baca.)
- [Ya92] M. Yannakakis, “On the Approximation of Maximum Satisfiability,” *J. of Algorithms*, Vol. 17, 1994, pp. 475-502.